*Please cite published version

How to Minimize Ontological Commitments: A Grounding-Reductive Approach

Abstract: Some revisionary ontologies are highly parsimonious: they posit far fewer entities than what we quantify over in ordinary discourse. The most radical examples are minimal ontologies, on which physical simples are the only things that exist. Highly parsimonious ontologies, and especially minimal ones, face the challenge of either accounting for the truth of our ordinary quantificational discourse, or paraphrasing such discourse away. Common strategies for addressing this challenge include classical reduction (by means of formal derivation and postulates), paraphrase nihilism, and a distinction between ontological and existence commitments. I argue, however, that these strategies are either implausible or fail to provide truth conditions consistent with minimal or parsimonious ontologies. I then discuss, defend, and suggest ways to strengthen an alternative framework for reduction, on which the sentences of reducing theories *ground* those of reducible theories. Relative to the other options for defending minimal ontology, a strengthened grounding-reductive approach can (in principle) provide more defensible truth conditions for minimal ontology, better preserve scientific realist intuitions, set a more attainable standard for reduction, and allow our existence commitments to be more responsive to empirical evidence and scientific expertise. As a result, I argue that minimal ontology becomes more defensible—though not certain—on a grounding-reductive framework. But even if minimal ontology were wrong, the grounding-reductive framework makes other parsimonious but non-minimal ontologies more plausible.

1. Minimal Ontology

Consider the following argument:

> **MIN:**
>
> (1) The only entities whose existence we're committed to are those required to make true the sentences of our best theory
>
> (2) The only entities required to make true the sentences of our best theory are physical simples
>
> (3) Therefore, the only entities whose existence we're committed to are physical simples

MIN-style arguments have been discussed in the literature for over a decade.[1] Such arguments converge on a *minimal ontology* (Cameron 2010), or an ontology that includes only physical simples (which have no proper parts). Given our best current physics, physical simples may include only subatomic particles (the various leptons, bosons, and quarks). So some minimal ontologists think these particles are all that exists. Alternatively, one might think the only things that exist are tropes. And on a minimal ontology, the only tropes that exist might be the determinate masses, charges and spins of the elementary particles—such particles being, for a certain kind of bundle theorist, nothing but collocated or concurrent tropes (Morganti 2009; Keinänen, Hakkarainen, & Keskinen 2018, 2019). Some minimal ontologists likewise posit nothing other than space-time points and perhaps sets of such points, for a so-called "pointillist supersubstantivalism" (Sider 2011, 2013; Dorr 2018). Even more radically, strong forms of ontic structural realism claim there are no objects at all.[2] (Weaker forms claim only that physical structure and relations are ontologically *prior* to entities like space-time points and subatomic particles; the latter might still exist) (Ladyman 2007; Ladyman and French 2011).[3] I won't argue in favor of one of these theories over the other. I take any one of these theories to count as a minimal ontology.

There are multiple motivations for minimal ontology. One might think, for example, that:

- Qualitative *ontological* parsimony is an important theoretical virtue, so we should commit to as few *types* of *entities* as possible. Minimal ontologies, especially pointil-

[1] See, for example, Heil (2012); Armstrong (2004); Cameron (2008, 2010); Sider (2011, 2013); Rettler (2016); and Dershowitz (2020). Of course, some of these arguments make less radical claims—that, for example, while our best *ontology* might include only physical simples, some composites might still *exist*, given a distinction between existence and ontology (Cameron 2010; Sider 2013; Rettler 2016). See section 2.3.

[2] Compare Chalmers (2009, 118), who suggests that the quantum wave function of the entire universe *could* provide the basis for an object-free view of reality, though he notes that this might leave "the character of the fundamental level of reality more obscure than the alternatives." Indeed, it may be difficult to see how such a view could allow for a sufficiently fine-grained characterization of local phenomena. Ney (2013) provides a more detailed framework for cashing out "wave function realism", though still at a rather high level of abstraction.

[3] Very roughly, OSR seeks to reconcile ontology with relativity and quantum mechanics by conceiving of reality as *fundamentally* a network of relations within a physical structure. But the more modest forms of ontic structural realism would imply only that objects lack haecceities or intrinsic properties, not that they don't exist at all.

list ones, would maximize qualitative parsimony. For there would only be physical simples.[4]

- Insofar as they eschew a parthood operator, and make do only with an operator for set membership, minimal ontologies may also maximize *ideological* or conceptual simplicity (Sider 2013; Brenner 2021).
- Fundamental physics is the best guide to what exists. Minimal ontologies are thus on the right track insofar as they needn't posit entities other than those required by fundamental physics.
- Mereological nihilism is true: no objects ever compose another object, so there is no such thing as a composite object, or an object with proper parts (Rosen and Dorr 2002; Dorr 2005; Cameron 2010; Sider 2011, 2013; Dershowitz 2020).[5] For the nihilist, then, physical simples are all that exists.

There are also some obvious sources of *resistance* to minimal ontology. For one thing, it may seem that common-sense intuition militates strongly against minimal ontologies. We often quantify over macroscopic objects without hesitation, and would think it odd not to do so. Of course, at least for those with a passing exposure to physics, there may *also* be an intuition that composites can be accounted for—ontologically, if not epistemically—through the interactions of subatomic particles, so that everything is really only physical simples, in the end.

In any event, insofar as minimal ontologies are unintuitive, their unintuitiveness might be dismissed as an artifact of a loose but convenient way of speaking. What's harder to dismiss is the following:

---

[4] The only alternative way to maximize parsimony would be a kind of *existence* monism, on which the world is the only thing that exists (e.g. Horgan and Potrc 2006). But this is perhaps even more counterintuitive than standard minimal ontologies, especially if one maintains that the world has no proper parts of its own. That is, existence monism commits us to the existence of worlds while eschewing commitment even to space-time points or fundamental particles. That may seem to get things backward. In any event, I won't be considering existence monism here. (Existence monism is distinct from *priority* monism, which holds that the world is more fundamental than its parts but that such parts have independent existence (Schaffer 2010)).

[5] Van Inwagen (1990) and Merricks (2001) famously argue that most of the entities we take to be composites don't exist, but that *some* composites do exist (which include organisms for van Inwagen (1990) and (roughly) persons for Merricks (2001)).

MAIN CHALLENGE: On minimal ontology, all our quantificational talk—even in our best scientific theories—would be false, aside from sentences quantifying over physical simples

The main challenge may seem to have intolerable implications. For it would (apparently) require us to accept that sentences quantifying over molecules and even atoms are all false. And while the minimal ontologist may not think that ordinary language is a reliable guide to ontology, it would be hard to be so dismissive of our best scientific theories. (Insofar as we should be wary about taking ordinary language as a guide to what exists, this would presumably be *because* science is a better guide).

There's also an analogous problem for ontologies that are parsimonious but not minimal. The problem is that parsimonious ontologies might still have us endorse an error theory about *most* of our ordinary quantificational talk, even if we admit *some* composites, such as molecules or basic chemical compounds. Such a parsimonious (but not fully minimal) ontology should also be a live option.

The most venerable way of dealing with the main challenge, and the spectre of an error theory, lies in classical reduction. In Quinean terms, we'd be committed to the existence of whichever entities $x_1...x_n$ are quantified over in the sentences $S_1...S_n$ of our best theories, *unless* $S_1...S_n$ could be reduced to sentences $S_1^*...S_n^*$ not quantifying over $x_1...x_n$—in which case we wouldn't need to commit to $x_1...x_n$ (Quine 1961, 103-104). If $S_1^*...S_n^*$ quantify only over simples, or only over microphysical entities, then we could rid ourselves of commitment to many ordinary things. And we might even hold that $S_1...S_n$ are true so long as they're fully reducible to $S_1^*...S_n^*$.

But as I explain later (in §2.4), classical reduction relies on a framework of formal derivation which has few examples in scientific practice, and has thus come to seem increasingly implausible. I argue, however, that other approaches to defending minimal ontology also fail. In particular, both paraphrase nihilism and a distinction between ontological and existence commitments fail to provide adequate *truth conditions* for a minimal ontology (see §2.1 and §2.2). Accordingly, in §3 and §4, I defend an alternative framework for reduction, on which the

sentences of reducing theories *ground* those of reducible ones (see Rosen 2010). Relative to the other options, such a grounding-reductive approach can (in principle) provide more defensible truth conditions for minimal ontology, stay more consistent with scientific realist intuitions, set a more attainable standard for reduction, and allow our existence commitments to be more responsive to empirical evidence and scientific expertise (see §4 and §5). As a result, I argue that minimal ontology becomes more plausible—though not certain—on the grounding-reductive framework. But even if minimal ontology doesn't obtain, I also argue that the grounding-reductive framework makes other parsimonious but non-minimal ontologies more plausible.

2. Some Approaches to Minimal Ontology

2.1 Paraphrase Nihilism

Now one approach to the main challenge is to admit that almost all our quantificational talk is false, but to maintain that very similar sentences are true. The main challenge would thus lose much of its bite. Let C-terms be terms purported (on typical interpretations) to refer to composites. Then instead of quantifying over Cs, we might say that some *simples* are arranged C-wise (Merricks 2001). We might even resort to plural quantification, and say that *there are* some simples arranged C-wise. So, for the paraphrase nihilist, even though it would be false that "there is a table," it may be true that "some simples are arranged table-wise."

Now paraphrase nihilism has been thoroughly discussed, and I can't provide an exhaustive critique of it. I just want to point out some of paraphrase nihilism's main deficiencies, so it can then be explained how my preferred option (in section 3) can improve upon them.

The first problem is that one needs an account of what it means for some simples to "be arranged" somehow. Presumably, to say that some simples $x_1...x_n$ are arranged C-wise is to say that each of $x_1...x_n$ is located relative to each other such that $x_1...x_n$ *would* count as composing a C, *if* there were indeed any facts about composition. Merricks (2001, 6) claims, for example, that to be arranged C-wise is (roughly) to bear the properties and relations upon which, if Cs existed, those Cs would supervene.

But this leads into another problem. For the nihilist claims that composites don't in fact exist. And yet the nihilist would be referring to composites to cash out the truth conditions for nihilist paraphrases. We'd be trying to paraphrase away terms—those purported to refer to composites—with paraphrases that employ *these same* terms as a kind of prefix, and with truth conditions that employ the terms themselves. At the very least, this seems like an incomplete paraphrase (Tallant 2014; Wilkins 2016).

There are thus two options. Either the references to composites in nihilist paraphrases are *reducible*, or they aren't. If they aren't, then it's not clear whether we really have good reason to resist allowing composites into our ontology. For if we *must* refer to composites (even in counterfactuals) to cash out the truth conditions for our talk of composites, that would suggest that composites are somehow part of reality. One might then worry, as Tallant (2014) does, that it's not clear why the truth conditions for nihilist paraphrases should not *themselves* count as criteria of composition. The nihilist, of course, could *stipulate* that such truth conditions are not criteria of composition, since composition never obtains (according to the nihilist). But is such a stipulation well-motivated? It might seem ad hoc at best. For if some reference to composites is necessary to cash out the truth conditions of nihilist paraphrases, the obvious approach is to abandon nihilism and admit that composition does obtain after all.

If, on the other hand, the references to composites are reducible in nihilist paraphrases, then there would be some presumably more *fundamental* truth conditions to which such references could be reduced. And it would be these truth conditions that the minimal ontologist should be interested in (Sider 2011, 2013). For talk of "simples arranged C-wise" would then just be a kind of shorthand or heuristic for the more fundamental truth conditions. And this brings out the larger point that I want to make: that the paraphrase nihilist still relies too much on natural language (or at least the ordinary, non-technical terms therein). The minimal ontologist, at least, wants to say that ordinary language is a highly unreliable guide to what exists, and we should look instead to fundamental physics for our ontology. To focus on providing ordinary-language truth conditions for our use of C-terms—especially in lieu of attempting a

more complete reduction—is thus to play on the home turf of those who already want to admit the existence of composites.

2.2 Ontology Versus Existence

Now another approach to minimal ontology focuses instead on criteria for what we're ontologically committed to—what we take to lie within our ontology. In particular:

> **TRUTHMAKER CRITERION:** For all x, we are ontologically committed to x iff x is among those entities necessary for the sentences of our best theories to be *made* true

Cameron (2008, 2010) makes the further claim that there's an important distinction between being committed to the *existence* of x and being *ontologically* committed to x (Cameron 2010; Rettler 2016).[6] On such a distinction, there would be a separate domain of *ontology*—distinct from mere existence—and ontology would be what's of most interest to metaphysics. Given the truthmaker criterion, the entities populating our ontology may only be a particular *subset* of the entities that exist—namely, that subset which includes all and only those entities required to serve as *truthmakers*. But since the only entities required as truthmakers would (according to Cameron) be physical simples, our ontology would include only physical simples. We would thereby get maximum (qualitative) *ontological* parsimony. After all, ontological parsimony is presumably about ontology!

The maneuver is tempting. But I think it runs into a dilemma. That is, either (1) there are at least some uses of "there is a C" (where 'C' purportedly refers to a composite, on typical interpretations) that should be taken as committing to *objective* facts about what *exists*, or (2) there are no such uses.[7] The dilemma is thus:

(A) If (1) is the case, then it's implausible that (mere) existence commitments don't compromise parsimony to *some* extent

(B) If (2) is the case, then we face largely the same problems as paraphrase nihilism

---

[6] One might also try to make a distinction between existence and fundamental existence do similar work (Sider 2011, 2013).

[7] Say that objective facts hold independently of context, convention, and attitude.

Regarding (A), suppose Cameron is right that our *ontology* is concerned only with the most fundamental entities, or the ones that serve as truthmakers. Then it's *not* implausible that positing *truthmaking* entities exacts a *greater* cost, in terms of decreased parsimony, than positing non-truthmaking entities that lie "outside" our ontology. What is implausible is that positing non-truthmaking entities exacts *no* cost—if those non-truthmaking entities do *in fact* exist. One can't minimize profligacy simply by designating some entities non-ontological, if those latter entities do in fact exist. For *objective* facts about the existence of *more* entities would mean there would be more entities in the *world*. I see no principled way to maintain that this wouldn't compromise parsimony, unless we radically change the meaning of "parsimony." If, then, composites do exist—even if they aren't in our ontology proper—this would undermine the minimal ontologist's motivation of maximizing qualitative parsimony.

If on, the other hand, the *non*-truthmaking entities (the composites, for Cameron) do not *in fact* exist, then something other than their existence would have to make true our ordinary ordinary-language sentences which seem to quantify over composites. What could this be? Well, according to Cameron, it would be the truthmakers themselves—the physical simples. The problem here is that this doesn't actually give us *truth conditions* for claims like "there is a table" (T_{table}). Indeed, it can't simply be a composite of simples that makes T_{table} true, for that would mean not only that composites exist, but that they're truthmakers—which Cameron explicitly rejects. But if it's not a composite that makes T_{table} true, just when is T_{table} true? Perhaps when certain simples are arranged table-wise! But then we face the same problems as paraphrase nihilism—namely, providing actual truth conditions for being arranged table-wise, while offering a principled rationale for why such conditions would *not* be criteria of composition in disguise.

Of course, it's not clear whether a distinction between ontological and existence commitments was really supposed to solve all the problems of paraphrase nihilism. The point, however, is that the problems remain. If we want to avoid the problems, we need to look elsewhere.

2.3 More Truthmaker Theories

One might embrace a truthmaker criterion without committing to such a controversial distinction between existence and real existence. Perhaps the minimal ontologist need only say that truthmaking is a kind of *grounding* between some worldly features (properties, states of affairs, perhaps even objects) and a proposition (Griffith 2014; O'conaill and Tahko 2016). (Grounding, as I'll explain in §3.1, is the typical way of cashing out non-causal and constitutive explanation in contemporary metaphysics). I'll argue in §3.4.3, however, that such an account of truthmaking as grounding is either unnecessary or inadequate for the purpose of reducing away our ontological commitments.

Of course, there are still other views of truthmaking. Schulte (2011, 2014), for example, provides a view of truthmaking intended to allow for explanatory reduction. Such a view draws on the notion of conceptual entailment; roughly, some x makes true p iff p is not analytic *and* a fundamental or base-level description of x's existence conceptually entails that p. As Schulte acknowledges, the notion of conceptual entailment is controversial. But whatever the merits of this notion of reductive explanation, it likely won't work for my purposes. Indeed, I doubt the notion of conceptual or analytic entailment alone is strong enough for *ridding* ourselves of ontological commitments. (As Chalmers and Jackson (1998) note, conceptual entailment alone needn't involve analysis or definition). Moreover, in line with contemporary realists, I take the fundamental issue in reductive explanation to be *ontological*. Conceptual entailment is *epistemic*, however. So whether there are conceptual entailments between the sentences of reducing and reducible theories may be an issue distinct from that of ontological reduction.

2.4 Classical Reduction

Reduction may be another option for answering the main challenge (see e.g. Sider 2011, 294-295). Indeed, I suggested as much in the introduction and at the end of 2.1. Now as classically conceived (e.g. by Nagel 1961), reduction is about the relation between the *sentences* of a (typically more fundamental) theory T_B and the sentences of a (typically less fundamental) theory T_R. Given some set of reduction postulates, a reduction of T_R to T_B occurs iff the sen-

tences of T_R are logically derivable from those of T_B. In order for such *classical* reduction to support minimal ontology, the sentences of our best theories quantifying over composites would have to be expressible in terms such that they could be formally derived from sentences quantifying over only simples. Namely:

REDUC:

(1) For all x, we need only commit to the existence of x iff x is quantified over in a sentence S of our best theories *and* S *cannot* be reduced to a sentence not quantifying over x

(2) For any sentence S of our best theories that quantifies over composites, S can be formally derived (via reduction postulates) from a sentence not quantifying over composites

(3) If a sentence S_R can be derived (via reduction postulates) from another sentence S_B, and S_B is true, then we can interpret S_R as true

(4) Therefore, we can interpret as true the sentences of our best theories quantifying over composites, while *not* committing to the existence of composites

Premise 1 of REDUC is based off of the classic Quinean criterion of ontological commitment (Quine 1948; 1961).[8] Premise (3) of REDUC may not be an uncontroversial metasemantic commitment, but I'm happy to assume it for now.

If REDUC were sound, we could answer the main challenge. The problem is that premise (2) of REDUC is highly speculative, as is the entire program of classical reduction. Let $S_{R,1...}S_{R,n}$ stand for the sentences of the *reduced* theory T_R, and let $S_{B,1...}S_{B,n}$ stand for the sentences of the base theory T_B—the theory *to which* the sentences of T_R are to be reduced. Then classical reduction would require *reduction postulates* or laws, such that T_R could only be reduced to T_B iff

[8] Indeed, on the Quinean criterion, we are ontologically committed to whichever entities are needed by the domain for the sentences of our best theories, regimented into first-order logic, to be evaluable as true, unless such sentences can be reduced to sentences not quantifying over such entities (Quine 1961; Schaffer 2008; Wetzel 2009, 28). If, for example, sentence S stipulates that ∃x x=a, then S is not true unless the entity denoted by 'a' exists, or unless S can be further reduced to a sentence which does not assign 'a' as the value of a variable (Schaffer 2008, 8). (The same goes for ∃X X=P, where 'P' is the name of a property, if we allow second-order quantification).

$S_{R,1...}S_{R,n}$ were expressible in terms such that there were some laws by which $S_{R,1...}S_{R,n}$ could be formally derived from $S_{B,1...}S_{B,n}$.

The problem here is that it's not clear there are such reduction postulates even for the reduction of chemistry to physics, much less the reduction of biology and the special sciences to physics (Hendry 2010a,b; Hettema 2012). At least, the current state of chemistry offers little prospect of a purely mathematical derivation of chemistry from physics. Even something like the structure of a molecule cannot be reduced via postulates to the terms of quantum mechanics (Hendry 2010b).

2.5 Ontological Reduction

In response, one might point to a distinction between classical reduction and *ontological* reduction that some propose (Hendry 2010a, 187; LePoidevin 2005). The former is a semantic and logical relationship between *sentences* of theories. The latter applies between entities, properties, or other *ontological* posits. There have been a variety of proposals for *ontological* reduction. An early example of ontological reduction—predating even Nagel (1961)—is given by Putnam and Oppenheim (1958), though this account would be unacceptable to nihilists given its appeals to mereological relations (see Ney 2013). Of course, most earlier reductivists still preferred to keep reduction largely on what Quine (1948) called the "semantic plane."

In recent years, accounts of ontological reduction have become more popular.[9] Such accounts typically appeal to non-causal or constitutive relations in addition to (or in lieu of) causal relations (LePoidevin 2005; Gillett 2007). One particularly rich account is given by Gillett (2007), who talks about non-causal "compositional" relations. The idea is that, very roughly, the mechanisms *constituted* by the causal *powers* of more fundamental (realizer) properties can account for the mechanisms *constituted* by the causal powers of less fundamental (real-

[9] This may reflect increasingly favorable attitudes towards metaphysics in general, and a retreat from the "linguistic turn" (Heil 2003). The grounding-reductive method that I discuss below may count as a prime example of such a metaphysical approach; but as I discuss, this approach is consistent with an empiricist epistemology.

ized) properties.[10] At least if we think that a property's theoretical role is largely exhausted by its causal powers, then it would be redundant to commit to the existence of both the higher-level and lower-level properties. So, applying a principle of ontological parsimony, we should commit to the existence of only the lower-level properties.

A problem with this kind of story is that it may seem *too* metaphysically committed. Individuating properties based on causal powers is not uncontroversial, and might be unpalatable even for some realists—especially those of a Humean variety who are skeptical about putting too much ontological weight on causation (and especially on the intrinsic causal powers of objects). Moreover, while mechanisms are frequently invoked in philosophy of science, it can be difficult to develop a consistent account of the notion of mechanism, or to precisify it formally. So there's reason to be cautious about assigning mechanism a central place in the metaphysics of reduction.

Now that's not to say that such a story is wrong. But it is to say that it's a species of a more general kind of metaphysical story, built around a notion of non-causal or constitutive explanation. My aim here is primarily to get clearer about this more general story, and how it could allow for minimizing ontological commitments. The general story might then be compatible with a variety of more specific accounts of ontological reduction.

3. The Grounding-Reductive Approach

3.1 Accordingly, we can take a step back and consider the basic intuitions behind reduction. There are different ways of framing such intuitions, but one way would be: what we observe of the world can be *explained* by facts about the most basic elements of reality, or physical simples. Consider again the sentence, "there is a table" (T_{table}). It seems quite intuitive, given what we know of physics, that T_{table} is *explained* by facts about the interactions of certain particles, particles which the paraphrase nihilist would say are "arranged table-wise."

[10] A favored example is that of neural mechanisms, in which the arrangement of protein sub-units account for the opening and closing of voltage-gated ion-channels, which account for the transmission of action potentials in neurons (Gillett 2007, 197; Craver and Darden 2001, 112-137).

Recently, then, there's been more focus on the nature of non-causal explanatory claims, as opposed to the search for particular reduction postulates. In contemporary metaphysics, non-causal or constitutive explanation is typically cashed out in terms of *grounding*. Grounding claims are expressed in ordinary language using phrases such as "in virtue of" or "because." And grounding, although itself hyperintensional, is typically taken to both imply and explain necessities (Rosen 2010; Fine 2012; Dasgupta 2014). So, for example, to say that *there is water* in virtue of the fact that *there is H_2O* is to say that the fact that there is H_2O both necessitates *and* (non-causally) explains the fact that there is water.

Now classical reduction seemed to set an (almost) unattainable standard for minimal ontology. Namely, the sentences of our best theories quantifying over composites would have to be expressible in terms in which they would be formally derivable, through reduction postulates, from sentences quantifying only over simples. Few examples of such reductions have been forthcoming, and so even the possibility of such reductions remains highly speculative. Grounding claims, on the other hand, can be (to some degree) established *empirically*, without the need for (purely) *formal* derivation of the reducible theory T_R from the reducing theory T_B. That is, the absence of such derivation does not itself prevent the (facts expressed by) the sentences of T_B from grounding those expressed by T_R, nor does it prevent there being sufficient empirical evidence to conclude that such grounding does obtain. This reflects the realist intuition that the limits of our *theorizing* aren't themselves a constraint on the way the *world* is.

What exactly counts as sufficient evidence for a grounding claim is, again, a question that is best left to the sciences themselves. It's fairly safe to say, for example, that we're far from being able to reduce something like phenomenal properties to microphysical ones, even eschewing the need for reduction postulates and formal derivation. On the other hand, we're likely much closer to being able to reduce certain macrophysical properties of chemical compounds to microstructural properties (Tahko 2015).

In any case, just because establishing a grounding claim doesn't require a formal derivation, that doesn't mean anything can be reduced to anything. On the contrary, the realist would hold that there are objective facts about which facts ground which other facts. So it's not the

existence of grounding facts that's up to the sciences—the facts would still be what they are whether or not they're discovered. The idea would rather be to defer *epistemically* to the relevant disciplines about whether there is enough evidence to *hypothesize* the existence of some particular grounding fact.

How, then, would the grounding-reductive approach work? Consider the following *stipulations*:

(*) Let $S_{R,1}...S_{R,n}$ be all the sentences of a theory T_R

(**) Let $S_{B,1}...S_{B,n}$ be all the sentences of a theory T_B

(***) If there are no other sentences $S_{B*1...}S_{B*n}$ that ground the sentences of $S_{B,1}...S_{B,n,}$ then (given (iii)), we say that T_B is a *base-level* theory.

(****) If T_R reduces to T_B, we can say that T_R is a *reducible* or *higher-level* theory, and T_B is a *reducing* or *lower-level* theory.

Now I'll sometimes talk about certain sentences of a base theory grounding the sentences of a higher-level theory. But I intend this to be a kind of shorthand, indicating that the facts or true propositions *expressed* by the base-level sentences ground the facts expressed by the higher-level sentences. As a first pass, then, one might suggest:

Grounding-Reductive Method (GROUND-REDUC):

$S_{R,1}...S_{R,n}$ are *reduced* to $S_{B,1}...S_{B,n}$ *iff* $S_{B,1}...S_{B,n}$ are true and $S_{B,1}...S_{B,n}$ *ground* the facts expressed by $S_{R,1}...S_{R,n}$ (see Rosen 2010, for example) (This is the *grounding-reductive* claim)

3.2 To see whether GROUND-REDUC succeeds, we need to consider what the *purpose* of a reduction is, and how *strong* a reduction needs to be to fulfill that purpose (Rosen 2010; Dorsey 2016). If the purpose of the reduction is merely to determine whether, and which, higher-level facts depend on lower-level ones, then something like GROUND-REDUC may be sufficient. If the purpose is instead to eliminate entities from our ontology—which is the goal of parsimonious ontologies—then GROUND-REDUC may fall prey to some obvious counterexam-

ples. Accordingly, I think GROUND-REDUC needs to be strengthened in order to provide a basis for *eliminative* reduction.

Indeed, grounding is typically taken to imply a necessary *conditional*: If A grounds B, then $\Box(A \supset B)$ (Sider 2013). So if all that reduction requires is a kind of metaphysical explanation, a *mere* grounding claim may suffice. A mere grounding claim, however, *isn't* generally held to imply a necessary *biconditional*.

Suppose, for example, that mental states are fully grounded by neuronal states. (Or that, for any mental state M and neuronal state P_N, the fact <there is some state M> is fully grounded by the fact <there is some neuronal state P_N>). At least assuming the *conditional* nature of grounding, this wouldn't mean that mental states couldn't *also* be fully grounded by silicon states, for example. (That is, it might be true *both* that <there is some silicon state P_S> fully grounds <there is some mental state M> *and* that <there is some neuronal state P_N> fully grounds <there is some mental state M>). On such assumptions, if we wanted to *remove* mental states from our ontology, it probably wouldn't be enough to replace mental states with neuronal states. For the same mental states could also be grounded (ex hypothesi) in silicon states. It would then seem as if mental states really were something *over and above* neuronal states, to use a common locution. So if we want a kind of reduction that warrants *eliminating* the reduced posits from our ontology, I think a mere grounding claim likely wouldn't suffice.

One could, of course, insist on so-called *full* grounding. That is, if A fully grounds B, then A is *sufficient* to ground B; nothing other than A is needed to ground B (Fine 2012). Even with full grounding, however, there is still the possibility of multiple realizability: if A fully grounds B, we can still infer only $\Box(A \supset B)$, not $\Box(A \leftrightarrow B)$.

3.3 To avoid such counterexamples, we might propose the following revision:

Revised Grounding-Reductive Method (GROUND-REDUC*):

$S_{R,1}...S_{R,n}$ are *reduced* to $S_{B,1}...S_{B,n}$ *iff* (i) $S_{B,1}...S_{B,n}$ are true, (ii) $S_{B,1}...S_{B,n}$ *ground* the facts expressed by $S_{R,1}...S_{R,n}$, *and* (iii) $\Box(S_{B,1}...S_{B,n} \leftrightarrow S_{R,1}...S_{R,n})$.

That is, GROUND-REDUC* adds a requirement of necessary co-extension between grounds and grounded. To minimize ontological commitments, we can then make the further claim:

> **EXIST-REDUC:** We needn't commit to the existence of any entities $x_1...x_n$ quantified over in the sentences of T_R, so long as the sentences of T_R can be reduced via GROUND-REDUC* to the sentences of a theory T_B which does not quantify over $x_1...x_n$

EXIST-REDUC is supported by the following argument:

(i) If $S_{B,1}...S_{B,n}$ provide truth conditions for, and fully ground, $S_{R,1}...S_{R,n}$, then we needn't commit to the entities $x_1...x_n$ quantified over in $S_{R,1}...S_{R,n}$

(ii) If $S_{R,1}...S_{R,n}$ are reducible via GROUND-REDUC* to $S_{B,1}...S_{B,n}$, then $S_{B,1}...S_{B,n}$ provide truth conditions for, and fully ground, $S_{R,1}...S_{R,n}$

(iii) Accordingly, if $S_{R,1}...S_{R,n}$ are reducible via GROUND-REDUC* to $S_{B,1}...S_{B,n}$, then we needn't commit to the entities $x_1...x_n$ quantified over in $S_{R,1}...S_{R,n}$

By way of illustration, suppose we had a complete list of *all* the possible physical grounds $P_1...P_n$ of a mental state M. Assuming further that M couldn't be ungrounded, the disjunction P_1 V P_2...V P_n would be necessarily co-extensive with M; and (ex hypothesi) P_1 V P_2...V P_n would also fully ground M. At least if we take parsimony seriously, there's a plausible case for eliminating M from our ontology. That's because of the combination of full grounding and necessary co-extension. Full grounding would secure the ontological and explanatory priority of P_B. And necessary coextension would allow complete *truth conditions* for the corresponding existential propositions: <there is some (property) M> would be true *iff* <there is some (property) P_1 V P_2...V P_n> is true.

Of course, if one finds disjunctiveness troublesome or too (ontologically) costly, one might require that the *reducing* (lower-level) properties be non-disjunctive. In that case, to continue with our prior example, the question would be whether there's some non-disjunctive lower-level property P_B that both fully grounds and is necessarily co-extensive with M. One way to *search for* P_B would be to determine whether there's some common property that *must*

be among the grounds of each of the disjuncts in P_1 V P_2...V P_n. (That is, whether there's a property P_B such that (P_B grounds P_1) & (P_B grounds P_2)...& (P_B grounds P_n)).[11] If so, then (by the transitivity of grounding) there would be some property P_B that grounds and is necessarily coextensive with M. So, given EXIST-REDUC, one could eliminate M in favor of P_B.

One might still be concerned, however, about whether grounding is the right tool for an eliminative reduction. Grounding theorists sometimes advocate a kind of tiered ontology, on which more fundamental (lower-level) entities such as physical simples ground less fundamental (higher-level) entities such as molecules. And such theorists typically take both the higher-level and lower-level entities to *exist* (e.g. Schaffer 2009). But grounding is not always considered a relation between *entities*. Grounding is often taken as a sentential operator, or an operator that applies to sentences or the propositions they express (Fine 2012; Dasgupta 2017). At least in my examples, I've taken the more pluralistic approach of allowing grounding to apply *both* between facts and between entities. But there's considerable precedent, especially in formal systems, for restricting grounding to facts or sentences alone. Such a restriction would ensure that grounding remains a type of *explanation*.

Indeed, GROUND REDUC* itself refers only to grounding between sentences or facts and propositions, and not between entities. Insofar as my examples deviate from this usage, they could be reframed to apply only to propositions. Accordingly, the quantificational talk in our higher-level theories might be taken to be about higher-level *predicates or concepts*, rather than about higher-level *entities or properties*. And as I'll discuss further below, I think we can have reason to retain concepts (especially outside of fundamental physics) that don't map one-to-one onto entities in our best ontology (see e.g. Gillett 2007; Sider 2011). The question would then be: which entities do we need to commit to in order to provide truth conditions for the sentences of our best higher-level theories?

[11] As a toy example, one might think of M as being a functionally individuated mental state, P_1 V P_2...V P_n as including all the possible physical realizers of M, and P_B as being a kind of abstract network that accurately characterizes all the possible physical realizers of M.

3.4.1 There are, of course, further objections to GROUND-REDUC*. Let's start with what may be the most formidable one. Namely, the *concept* of a higher-level property like M might play an inferential role not fulfillable by the *concept* of a lower-level property like P_B, even if M were reducible to P_B (by GROUND-REDUC*). Together with the assumption that inferential role should be a guide to ontology, one might then resist eliminating M from our ontology.

Now I wouldn't deny that there *could* be reason to retain the *concept* of M, even if it were reducible to P_B by GROUND-REDUC*. But I think we can have reason to retain *concepts* (especially outside of fundamental physics) that don't map one-to-one onto *entities* in our best ontology (see e.g. Gillett 2007; Sider 2011). Indeed, from a realist perspective, there should be a clean separation of the ontological from the epistemic and the conceptual. So it's not clear the realist should take inferential role to be a guide to ontology.

This last intuition—about the distinction between the ontological and the conceptual—may be the simplest and most forceful response the realist can make on behalf of GROUND-REDUC*. But if one is still troubled by whether inferential role gets in the way of the grounding-reductive approach, one might make a concession to the anti-reductivist. That is, one could say:

> **INFERENTIAL CONCESSION**: An entity x exists iff (i) x is quantified over in a sentence of our best theories that's *irreducible* according to GROUND-REDUC* *or* (ii) the *concept* of x plays an indispensable role in inference.

Natural kinds might be prime candidates for the second disjunct of the inferential concession. One of the main arguments for positing natural kinds is that they permit a very wide range of inductive inferences (Bird 2018).[12] Nonetheless, kinds—or rather, kind *concepts* – may only be inferentially "indispensable" because of the limits of human cognition or representation. And the minimal ontologist can *consistently* claim that a qualitatively parsimonious picture eschewing kinds best captures the ontological structure of the world, or carves closest to

[12] One might think in particular of species, which may well be individuated according to evolutionary history. Of course, even in the case of species, genetic and structural factors may also play some role in individuation (Devitt 2008, 2018; LaPorte 2017).

nature's joints, while allowing that it would be impractical to eliminate kind *concepts* from our theories given their role in inference.

Indeed, I don't think a *stronger* realist—one who places great weight on the distinction between the ontological versus the conceptual or representational—has to make the inferential concession. Weaker realists might be tempted towards such a concession, admittedly. Of course, if one makes this concession *and* it turns out that our theories do employ higher-level concepts which are indispensable for inference, then one would have to retreat from minimal ontology *simpliciter* (in one of the senses given in §1). But one might still use the grounding-reductive method to significantly *decrease* our ontological commitments. That is, we could eliminate any entities that *weren't* inferentially indispensable but that could be reduced by the grounding-reductive method.

3.4.2 Secondly, one might think that GROUND-REDUC* falls afoul of the irreflexivity of grounding. For GROUND-REDUC* requires necessary coextension between grounds and grounded. But if p is *identical* to q, one might think that p can't also ground q, because grounding (like explanation itself) is an *irreflexive* relation. So if there were an *identity* between the (facts expressed by) the sentences of higher-level and lower-level sentences, perhaps the lower-level sentences couldn't also ground the higher-level ones without violating irreflexivity (Dorsey 2016).

Nonetheless, all GROUND-REDUC* requires is that there be necessary coextension (and full grounding) between the facts expressed by the sentences of the reducible and reducing theories. And, *pace* Hume and Lewis, I don't think necessary coextension is a good criterion of *identity*. The classic counter-example here is that of equilaterality and equiangularity. The two properties are necessarily co-extensive (along with the corresponding propositions that <x is equiangular> and that <x is equilateral), but intuitively, the properties are clearly different. Angles and lengths are two quite different things (see Sober 1982; Eddon 2011). Accordingly, if there isn't an identity between the facts expressed by the reducible and reducing theories, then the irreflexivity of grounding isn't relevant here.

But even if one thinks there must be an *identity* between the reducible and the reducing (or that necessary coextension *does* imply identity), there's a response available. Namely, one can replace the grounding operator in GROUND-REDUC* with another operator ($\equiv_{def}$) which *implies* an identity, but is itself *irreflexive* (see Correia 2017).[13] $\equiv_{def}$ is how Correia (2017) cashes out real definition. Given such irreflexivity, one could *not* say that "to be G $\equiv_{def}$ to be G," although one obviously could say that "to be G $\equiv$ to be G" (where $\equiv$ is an operator for generalised identity). The motivation for the rules of $\equiv_{def}$ is, roughly, that certain concepts and facts appear to ground or have ontological priority over other facts, and yet still imply an identity. Consider the stock example: to be water is to be (a certain microstructure of) H_2O molecules. Since the $\equiv_{def}$ operator allows such a priority of the right-hand side over the left without violating irreflexivity, it might seem well-suited to the task of reduction, if one thought reduction needed to imply identity as well.[14]

3.4.3 Finally, one might ask whether a truthmaker criterion of existence commitment, together with a grounding-based account of truthmaking, might not offer a better approach to reduction (see §2.2 and §2.3). (Recall that truthmaking-as-grounding holds that it's grounding relations between worldly features and propositions that are best suited to explain the truth of true propositions). But I think the view of truthmaking-as-grounding—even if it were true—might add unnecessary complexity to a *reductionist* story. It's not clear, in particular, that the relation between *propositions* and ontic features (properties, relations, objects) is best captured by grounding. One can hold that the facts expressed by the sentences of the reducing theory *ground* those of the reducible theory, while remaining *agnostic* about whether and what kind of truthmaking relation is necessary to account for the truth of the relevant facts. Indeed, one might think that the truthmaking relation isn't about grounding, but about how the semantic properties of facts relate to entities (see e.g. Audi 2020).

---

[13] Very roughly, given a version of Leibniz's law, this creates a paradox, which Correia (2017, 56) resolves by holding that $\equiv_{def}$ is *opaque* on both sides. Given opaqueness, we can then accept certain violations of Leibniz's law.
[14] Now, as Correia (2017, 58) also discusses, the irreflexivity of $\equiv_{def}$ is more plausible with representational as opposed to "worldly" grounding. Representational grounding, very roughly, concerns propositions and concepts, whereas *worldly* grounding involves entities "in the world," such as concrete states of affairs and properties.

But even setting aside such concerns, if truthmaking-as-grounding is supposed to allow for eliminative reduction, it may fall prey to the same counterexamples as the initial GROUND-REDUC (as opposed to GROUND-REDUC*). Suppose, that is, that some entities $x_1...x_n$ ground the truth of some proposition p, where p is about some mental state M. Can one eliminate M from one's ontology in favor of $x_1...x_n$? Well, not necessarily; for p (the proposition about M) might also be grounded by $y_1...y_n$, given the conditional (rather than biconditional) nature of grounding. To resist such counterexamples, the only obvious recourse for the truthmaker theorist would be to adopt something like Cameron's distinction between ontology and existence. But I've already explained (in §2.2) why this distinction is problematic at best.

4. Reduction and Minimal Ontology

4.1 How then does the framework of GROUND-REDUC* support minimal ontology? What would have to be the case for us to be warranted in committing only to the existence of physical simples? If one accepts GROUND-REDUC*, the obvious answer may be:

> **Minimal Existence Commitments (MIN-EXIST)**: We should commit *exclusively* to the existence of physical simples, and *not* to the existence of any composites *iff* the sentences $S_{R,1...}S_{R,n}$ of all our best theories employing C-terms are reducible—by the criteria of GROUNDING-REDUC*—to sentences of a base theory T_B not employing any C-terms.

I've argued above that paraphrase nihilism alone won't suffice for minimal ontology. Despite their professed commitment to an ontology of simples, the paraphrase nihilists haven't been able to show how we can avoid admitting composites into our ontology. And the *classical* reduction of all scientific theory to fundamental physics remains implausible. Accordingly, the minimal ontologist needs an alternative framework, in order to show that a minimal ontology even remains viable *in principle*.

GROUND-REDUC* provides that framework by furnishing the *truth conditions* for sentences $S_{R,1...}S_{R,n}$ employing C-terms. And if our base theory doesn't include any terms for composites, we thereby avoid the paraphrase nihilist's conundrum of giving truth conditions for $S_{R,1...}S_{R,n}$ that employ C-terms (even if they don't explicitly quantify over Cs).

MIN-EXIST might seem like a demanding requirement, but the notion of reduction in GROUND-REDUC* may make MIN-EXIST more plausible. Indeed, GROUND-REDUC* does not require reduction postulates or bridge laws, nor does it require *hyperintensional* identity of the facts expressed by the sentences of the reducible and reducing theories (although it does require necessary *coextension*). In any case, as I've argued above, one can discover grounding relations at least partly on the basis of empirical evidence, without the need for formal derivation.[15] This itself makes MIN-EXIST more plausible. For facts about the interactions of simples really might account for all the facts expressed by (the sentences of) our best higher-level theories. At least, the realist thinks there could be objective facts about whether this is so, and the minimal ontologist thinks it is so. Yet even if it were so, the limits of our *theorizing*—either our own cognitive limits, or the representational limits of the theoretical languages themselves—still might not allow for a complete formal derivation of the sentences of higher-level theories from those of lower-level ones.

In any case, I think we can then argue for minimal ontology itself as follows:

Composite Reduction (COMP-REDUC)

- Let $S_{B,1}...S_{B,n}$ be all the sentences of a base-level theory T_B, where T_B quantifies only over simples and includes no terms interpretable as referring to composites
- Let $T_{R,1}...T_{R,n}$ be all our best higher-level theories, and let $S_{R,1}...S_{R,n}$ be all the sentences of $T_{R,1}...T_{R,n}$ that quantify over $x_{R,1}...x_{R,n}$, where $x_{R,1}...x_{R,n}$ are all the entities in $T_{R,1}...T_{R,n}$ that are referred to by terms typically interpreted as referring to composites. Then:

[15] Now one might ask what the *criterion* is for determining when the sentences of the reducing (lower-level) theories ground those of the reducible (higher-level) theories. I think, however, that this question is best understood epistemically. That is, how can *we* tell when some sentences of a lower-level theory ground those of a higher-level one? As I've suggested in 3.1, however, I would leave such questions largely to the sciences themselves.

(1) We needn't commit to the existence of $x_{R,1}...x_{R,n}$ if (A) we can reduce (by the grounding-reductive method) sentences $S_{R,1}...S_{R,n}$ to sentences $S_{B,1}...S_{B,n}$, *and* (B) $S_{B,1}...S_{B,n}$ are true

(2) $S_{R,1}...S_{R,n}$ can be reduced (by the grounding-reductive method) to $S_{B,1...}S_{B,n}$

(3) $S_{B,1...}S_{B,n}$ are true

(4) Therefore we needn't commit to the existence of objects other than simples

Premise (1) I've already defended in §3.2 Premise (4) follows from the fact that if premises (1)-(3) are true, then all our best theories are reducible to T_B, and T_B includes *no* terms interpretable as referring to composites (and no parthood operator). Indeed, to each theory will correspond certain distinctive terms. For our higher-level theories, such as chemistry and biology, these terms may well include C-terms typically interpreted as referring to composites. T_B, on the other hand, would draw on the ideology (concepts and expressions) and ontology of fundamental physics and physical geometry, which don't include composites (Cameron 2010; Sider 2011, 2013). The idea would be that we can rid ourselves of commitment to composites by reducing all our best higher-level theories to T_B, according to the grounding-reductive method.

Premise (2) of COMP-REDUC, as I've suggested above, is a matter of whether the intuitions motivating minimal ontology are true. The question is whether facts about physical simples can explain all the facts expressed by our higher-level theories—and that doesn't seem to me to be a question that can be answered by metaphysics alone. Indeed, I've argued above that such grounding claims are best left for the relevant sciences to refute or support. But on the grounding-reductive approach, as opposed to classical reduction, it may become more plausible that a reduction of all our best higher-level theories to T_B is in fact possible (see §3.1). For the grounding-reductive approach would allow reduction to be done at least partly on the basis of empirical methods, without the need for a complete formal derivation.

4.2 One might also wonder how the sentences of our reducible theories—sentences employing C-terms—could be true at all if, as the minimal ontologist claims, composites don't ex-

ist. (Say that *C-terms* ("table," "chair," "molecule") are terms purported (on typical interpretations) to refer to composites). What other than composites would C-terms refer to? Here, however, there are *at least* two options for the minimal ontologist. One option would be to say that C-terms do refer, but that they don't refer to composites. Rather, they refer to something else, such as pluralities or sets of physical simples with no mereological structure (see section 5 below). So when we say, "There is a C," we wouldn't be in the business of committing to *composites*.

If, however, one denies that sentences quantifying over composites can be true, one might reframe such sentences according to the paraphrase nihilist formula. One might replace "there are tables," for example, with "some simples are arranged table-wise," or "there are some simples arranged table-wise." Such nihilist paraphrases might still be true, even if they need to be *reducible* to sentences of base-level theories in order to avoid commitment to composites. So it might be *these* paraphrased sentences that are reduced according to the grounding-reductive method. Such paraphrases, then, would just be a sort of heuristic or *façon de parler* for the base-level sentences to which they can be reduced.

One might still wonder what our base theory T_B looks like. After all, if T_B doesn't include any terms that might be associated with composites, what does it include? The best option here would likely resemble a *spatial* logic or language (Tarski 1959; Sider 2011).[16] Spatial logics are designed to capture logical and formal relations between geometrical structures (points, lines, surfaces, etc.). Accordingly, spatial logics may incorporate predicate logic, plus some *additional* operators for describing relations between points or regions. As Sider (2011, 2013) suggests, such spatial languages may also be expanded to represent physical geometry—that is, structures which exist in the world. Insofar as one wants such a language to apply to the world, and not only to abstract geometric structures, one can incorporate certain operators from physics, such as "is more massive than," into the language (Sider 2011, 294). Indeed,

[16] Sider (2011, 2013) has already discussed using spatial logics to reduce higher-level theories, though he hasn't provided the framework for reduction that I've defended in the previous section.

the atoms of such a language would not be propositions, but *points*. And such points may be taken to represent simples.

Still, COMP-REDUC, if sound, purports only to show that we *needn't* commit to the existence of composites. Why *shouldn't* we so commit? Well, I've already mentioned some of the motivations for minimal ontology: chiefly, qualitative parsimony, and the intuition that fundamental physics can explain or ground all the facts expressed by our higher-level theories. Sider (2013) emphasizes another possible motivation. Namely, a theory eschewing commitment to composites may be *ideologically* simpler—that is, it can make do with fewer primitive operators or expressions, by forgoing a parthood operator. So for those attracted to minimal ontology, there is motivation enough to eschew composites. The question is whether one can get away with doing so. I've argued one can, *if* COMP-REDUC is sound.

5. Minimizing Ontological Commitments

Now the grounding-reductive framework won't vindicate minimal ontology by itself, though I've argued it makes minimal ontology more plausible through its eschewal of reduction postulates. (That is, COMP-REDUC may be unsound even if the grounding-reductive framework is correct).[17] Indeed the framework doesn't purport to be sufficient for establishing minimal ontology. But this may actually offer certain advantages. For the result allows a more modest and empirical approach to ontology.

The discourse on nihilism typically focuses on whether composition obtains, and if so, which criterion of composition is correct. If GROUND-REDUC is correct, however, we have some finer-grained ontological options. It's not implausible, for example, that while all our higher-order theories can't be reduced to the terms of basic physical geometry, they could nonetheless be reduced to theories employing only microphysical properties. In that case, we would commit to the existence of both physical simples and whichever microphysical entities couldn't be reduced to simples. (Or rather, whatever wasn't quantified over in *sentences* that could be so reduced). On the other hand, it's also not implausible that our *only* irreducible commitments

[17] That is, the falsity of premise 2 of COMP-REDUC is consistent with all the criteria for GROUND-REDUC* being true. It's only premise 1 of COMP-REDUC that would require the criteria of GROUND-REDUC* to be true. See §3.1.

might be to phenomenal properties or neuropsychological states, in which case we'd commit only to the existence of these *plus* physical simples.

Thus we get some modesty. There may be sufficient motivation to take minimal ontology seriously—given its greater ontological parsimony, ideological simplicity, continuity with fundamental physics, etc. For all that, however, the thesis that there are only simples could still be false—unless one took maximal parsimony to be *necessarily* truth-conducive. But even if the thesis of minimal ontology were false, that wouldn't mean we need admit all the entities quantified over in ordinary language. Given the grounding-reductive framework, we still might commit to only a select group of entities, as I've suggested. And we need not defend such a commitment as following from a particular criterion of composition, selected based *only* on consideration of theoretical costs and benefits.[18] Which entities we commit to would be largely an empirical matter, determined based on the evidence gathered by the respective sciences about which claims of which theories do or don't ground which claims of which other theories.

Accordingly, the grounding-reductive framework allows us to combine a realist approach to ontology with an empiricist approach to epistemology. We can defer to the sciences about what does exist, on the basis of a fairly simple framework for reduction. If the minimal ontologist's intuitions are correct, we should be able to provide truth conditions in the terms of a base theory quantifying only over simples, without the need for reduction postulates. On the other hand, if it turns out that facts about simples *alone* aren't sufficient to account for the claims of our best theories, the framework offers a principled basis for an ontology that reflects this, without the need to justify some principle of composition *a priori*. We can thus minimize our ontological commitments to the degree allowed by our best science.

References

Armstrong, D. M. (2004). Truth and truthmakers. Cambridge: Cambridge University Press.

[18] Indeed, there are often counterbalancing theoretical costs and benefits involved in selecting such criteria. This complicates the attempt to select such a criterion without bringing empirical considerations to bear, even apart from epistemic concerns about whether theoretical virtues are truth-conducive (see Bennett 2009).

Audi, P. (2020). Why truthmaking is not a case of grounding. *Philosophy and Phenomenological Research*, *101*(3), 567-590.
Bennett, K. (2009). Composition, colocation, and metaontology. *Metametaphysics: New essays on the foundations of ontology*, 38-76.
Bird, A. (2018). The metaphysics of natural kinds. *Synthese*, *195*(4), 1397-1426.
Brenner, A. (2021). Mereology and ideology. *Synthese*, *198*(8), 7431-7448.
Cameron, R. (2008). Truthmakers and ontological commitment: Or how to deal with complex objects and mathematical ontology without getting into trouble. *Philosophical Studies*, *140*(1), 1–18.
Cameron, R. (2010). How to have a radically minimal ontology. *Philosophical studies*, *151*(2), 249-264.
Chalmers, D. J., & Jackson, F. (2001). Conceptual analysis and reductive explanation. *The Philosophical Review*, *110*(3), 315-360.
Craver, C. F., & Darden, L. (2001). Discovering mechanisms in neurobiology: The case of spatial memory.
Dasgupta, S. (2014). The possibility of physicalism. The Journal of Philosophy, 111(9/10), 557-592.
Dasgupta, S. (2017). Constitutive explanation. *Philosophical Issues*, *27*(1), 74-97.
Devitt, M. (2008). Resurrecting biological essentialism. *Philosophy of Science*, *75*(3), 344-382.
Devitt, M. (2018). Individual essentialism in biology. *Biology & Philosophy*, *33*(5), 1-22.
Dorr, C. (2005). What we disagree about when we disagree about ontology. *Fictionalism in metaphysics*, *234*, 86.
Dorr, C. (2018). A Way of Getting Rid of Things: Higher-order Languages, Priorian Nominalism, and Nihilism. MS.
Dorsey, J. E. (2016). On the grounding-reduction link. *American Philosophical Quarterly*, 411-422.
Dershowitz, N. (2020). Nihilism, but not necessarily. *Erkenntnis*, 1-16.
Eddon, M. (2011). Intrinsicality and hyperintensionality. *Philosophy and Phenomenological Research*, *82*(2), 314-336.
Fine, K. (2012). Guide to ground. *Metaphysical grounding: Understanding the structure of reality*, 37-80.
Griffith, A. M. (2014). Truthmaking and grounding. *Inquiry*, *57*(2), 196-215.
Gillett, C. (2007). Understanding the new reductionism: The metaphysics of science and compositional reduction. *The Journal of Philosophy*, *104*(4), 193-216.
Heil, J.. (2003). *From an ontological point of view*. Oxford: Oxford University Press.
Heil, J. (2012). *The universe as we find it*. Oxford: Oxford University Press.
Hendry, R.F. (2010). Emergence vs. reduction in chemistry. *Emergence in mind* (2010): 205-221.
Hendry, R.F. (2010). Ontological reduction and molecular structure. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *41*(2), 183-191.
Hettema, H. (2012). *Reducing chemistry to physics: Limits, models, consequences*. Groningen: Rijksuniversiteit Groningen.
Horgan, T., & M. Potrč. (2006). Abundant truth in an austere world. In Greenough, P. & M. P. Lynch (Eds.), *Truth and realism: New essays* (pp. 137-167). Oxford: Oxford University Press.
Keinänen, M., J. Hakkarainen, & A. Keskinen. (2018). Kinds of tropes without kinds. *Dialectica*, *72*(4), 571-596.
Keinänen, M., A. Keskinen, & J. Hakkarainen. (2019). Quantity tropes and internal relations. *Erkenntnis 84*(3), 519-534.
LaPorte, J. (2017). Modern essentialism for species and its animadversions. In *The Routledge handbook of evolution and philosophy* (pp. 182-193). Routledge.
LePoidevin, R. (2005). Missing elements and missing premises: a combinatorial argument for the ontological reduction of chemistry. *The British journal for the philosophy of science*, *56*(1), 117-134.
Morganti, M. (2009). Tropes and physics. *Grazer Philosophische Studien*, *78*(1), 185-205.
Nagel, E. (1961). *The Structure of Science*: *Problems in the Logic of Scientific Explanation*. Harcourt, Brace & World.

Ney, A. (2013). Ontological reduction and the wave function ontology. In Ney, A., & Albert, D. Z. (Eds.), *The wave function: Essays on the metaphysics of quantum mechanics*. Oxford: Oxford University Press.
Oppenheim, P., & Putnam, H. (1958). Unity of science as a working hypothesis.
O'conaill, D., & Tahko, T. E. (2016). Minimal truthmakers. *Pacific Philosophical Quarterly*, *97*(2), 228-244.
Quine, W.V.O. (1948). On what there is. *The review of metaphysics*, 21-38.
Quine, W.V.O. (1961). *From a logical point of view: Nine logico-philosophical essays*. Cambridge: Harvard University Press.
Rosen, G., & Dorr, C. (2002). Composition as a Fiction. *The Blackwell guide to metaphysics*, 151-174.
Rosen, G. (2010). Metaphysical dependence: Grounding and reduction. *Modality: Metaphysics, logic, and epistemology*, *109*, 109.
Rettler, B. (2016). The general truthmaker view of ontological commitment. *Philosophical Studies*, *173*(5), 1405-1425.
Schaffer, J. (2008). Truthmaker commitments. *Philosophical Studies 141*(1), 7-19.
Schaffer, Jonathan (2009). On what grounds what. In David Manley, David J. Chalmers & Ryan Wasserman (Eds.), *Metametaphysics: New Essays on the Foundations of Ontology* (pp. 347-383). Oxford University Press..
Schaffer, J. (2010). Monism: The priority of the whole. *The Philosophical Review*, *119*(1), 31-76.
Schulte, P. (2011). Truthmakers: a tale of two explanatory projects. *Synthese*, *181*(3), 413-431.
Schulte, P. (2014). Can truthmaker theorists claim ontological free lunches? *European Journal of Philosophy*, *22*(2), 249-268.
Sider, T. (2011). *Writing the Book of the World*. Oxford: Oxford University Press.
Sider, T. (2013). Against parthood. *Oxford studies in metaphysics* 8, 237-293.
Sober, E. (1982). Why logically equivalent predicates may pick out different properties. *American Philosophical Quarterly*, *19*(2), 183-189.
Tallant, J. (2014). Against mereological nihilism. *Synthese*, *191*(7), 1511-1527.
Tahko, T. E. (2015). Natural kind essentialism revisited. *Mind*, *124*(495), 795-822.
Wetzel, L. (2009). *Types and tokens: On abstract objects*. Cambridge: MIT Press
Wilkins, S. M. (2016). Why paraphrase nihilism fails. *Synthese*, *193*(8), 2619-2632.